

Identification of biologically relevant subtypes via preweighted sparse clustering

SHEILA GAYNOR, ERIC BAIR*

*Departments of Biostatistics and Endodontics, The University of North Carolina in Chapel Hill,
Chapel Hill, NC*

ebair@email.unc.edu

SUMMARY

Cluster analysis methods are used to identify homogeneous subgroups in a data set. Frequently one applies cluster analysis in order to identify biologically interesting subgroups. In particular, one may wish to identify subgroups that are associated with a particular outcome of interest. Conventional clustering methods often fail to identify such subgroups, particularly when there are a large number of high-variance features in the data set. Conventional methods may identify clusters associated with these high-variance features when one wishes to obtain secondary clusters that are more interesting biologically or more strongly associated with a particular outcome of interest. We describe a modification of the sparse clustering method of [Witten and Tibshirani \(2010\)](#) that can be used to identify such secondary clusters or clusters associated with an outcome of interest. We show that this method can correctly identify such clusters of interest in several simulation scenarios. The method is also applied to a large case-control study of TMD and a breast cancer microarray data set.

*To whom correspondence should be addressed.

Key words: Cancer; Cluster analysis; High-dimensional data; K-means clustering; Temporomandibular disorders.

1. INTRODUCTION

In biomedical applications, cluster analysis is frequently used to identify homogeneous subgroups in a data set that are associated with some outcome of interest. For example, in microarray studies of cancer, a common objective is to identify cancer subtypes that are predictive of the prognosis (survival time) of cancer patients ([Bhattacharjee and others, 2001](#); [Sorlie and others, 2001](#); [van 't Veer and others, 2002](#); [Rosenwald and others, 2002](#); [Lapointe and others, 2004](#); [Bullinger and others, 2004](#)). In studies of chronic pain conditions, such as fibromyalgia or temporomandibular disorders (TMD), one may wish to develop a more precise case definition for the condition of interest by identifying subgroups of patients with similar clinical characteristics ([Jamison and others, 1988](#); [Bruehl and others, 2002](#); [Davis and others, 2003](#); [Hastie and others, 2005](#)). However, conventional clustering methods (such as k-means clustering and hierarchical clustering) may produce unsatisfactory results when applied to these types of problems.

Identification of biologically relevant clusters in complex data sets presents several challenges. It is common for the biologically relevant clusters to differ with respect to only a subset of the features. This is particularly true in genetic studies, where the majority of the genes are not associated with the outcome of interest. Moreover, it is possible that some other subset of the features form clusters that are not associated with the outcome of interest. In genetic studies, given that genes work in pathways, genes in the same pathway are likely to form clusters even if the pathway is not associated with the biological outcome of interest.

As a motivating example, consider the (artificial) data set represented in Figure 1. Observe that there are two sets of clusters in this data set: features 1-50 form one set of clusters, and features 51-250 form a separate set of clusters. Also, note that the difference between the cluster

means is much greater for the clusters formed by features 51-250 than it is for the clusters formed by features 1-50. Thus, when conventional clustering methods are applied to this data set, they will most likely identify the clusters corresponding to features 51-250. However, if observations 1-100 are controls and observations 101-200 are cases, then we would be interested in the clusters corresponding to features 1-50, which would not be identified by most existing clustering methods. See [Nowak and Tibshirani \(2008\)](#) for a more detailed discussion of this problem.

A number of methods exist for clustering data sets when the clusters differ with respect to only a subset of the features ([Ghosh and Chinnaiyan, 2002](#); [Friedman and Meulman, 2004](#); [Bair and Tibshirani, 2004](#); [Raftery and Dean, 2006](#); [Pan and Shen, 2007](#); [Koestler *and others*, 2010](#); [Witten and Tibshirani, 2010](#)). In particular, the method of [Nowak and Tibshirani \(2008\)](#) is designed specifically for the situation described in [Figure 1](#). However, many of these methods are computationally intensive, and their running times may be prohibitive when applied to high-dimensional data sets. More importantly, with a few exceptions, these methods do not consider an outcome variable or any other biological information that could help identify the clusters of interest. In other words, if these methods are applied to a data set similar to [Figure 1](#), they are likely to produce clusters that are not related to the outcome of interest.

A wide variety of methods have been proposed for situations where an outcome variable is available and one wishes to identify clusters that are associated with the outcome variable ([Basu *and others*, 2002](#); [Bair and Tibshirani, 2004](#); [Qu and Xu, 2004](#); [Handl and Knowles, 2006](#); [Tari *and others*, 2009](#); [Koestler *and others*, 2010](#)). The majority of these methods assume that the true cluster assignments are known for some subset of the observations. However, in many situations, the outcome variable is a “noisy surrogate” ([Bair and Tibshirani, 2004](#); [Bair *and others*, 2006](#)) for the true clusters. In other words, the outcome variable provides some information about the clusters of interest, but the true cluster assignments are still unknown for all observations. An (artificial) example of this situation is shown in [Figure 2](#). In this example, the mean of the

outcome variable for observations in cluster 2 is higher than the mean of the outcome variable for observations in cluster 1. However, there is considerable overlap in the distributions. Thus, higher values of the outcome variable increase the likelihood that an observation belongs to cluster 2, but any classifier that attempts to predict the cluster based on the outcome variable will have a high error rate.

This situation where the observed outcome variable is a noisy surrogate variable for underlying clusters is very common in real-world problems. For example, in genetic studies of cancer, it is believed that there are underlying subtypes of cancer with different genetic aberrations, and some subtypes may be more responsive to treatment ([Rosenwald and others, 2002](#); [Bullinger and others, 2004](#); [Bair and Tibshirani, 2004](#)). These subtypes cannot be observed directly, but a surrogate variable (such as the patient’s survival) time may be available. Patients with a “low-risk” tumor type will have greater survival than patients with a “high-risk” tumor type. The goal of the study is to identify these underlying tumor subtypes using both the patient’s survival time and gene expression data. Similarly, the diagnosis of chronic pain conditions is based entirely on the patient’s subjective pain report. For example, TMD is diagnosed based on a patient’s self-reported pain levels after palpation at certain locations on the jaw ([Dworkin and LeResche, 1992](#)). Other pain conditions, such as dysmenorrhea, have no clearly defined case definition, and are diagnosed when a patient complains of pain that interferes with daily living ([Sanfilippo and Erb, 2008](#)). Since the diagnosis of these conditions is necessarily subjective, it is likely that many “cases” with a chronic pain condition are more similar to controls without chronic pain and that many “controls” have phenotypic characteristics similar to chronic pain patients ([Diatchenko and others, 2006](#)). An important goal in studies of chronic pain is to identify clusters of patients with similar symptoms even if the clusters do not correspond perfectly with the existing definitions of a specific painful condition.

Despite the fact that this situation is common in real-world problems, there are relatively

few clustering methods that are applicable for this type of problem. [Bair and Tibshirani \(2004\)](#) propose a method that they called “supervised clustering.” Supervised clustering performs conventional k-means clustering or hierarchical clustering using only a subset of the features. The features are selected by identifying the features that have the strongest univariate association with the outcome variable. For example, if the outcome is dichotomous, one would calculate a t-statistic for each feature to test the null hypothesis of no association between the feature and the outcome and then perform clustering using only the features with the largest (absolute) t-statistics. [Koestler and others \(2010\)](#) propose a method called “semi-supervised recursively partitioned mixture models” (or “semi-supervised RPMM”). This method is similar to the supervised clustering method of [Bair and Tibshirani \(2004\)](#) in that one first calculates a score for each feature (such a t-statistic) that measures the association between that feature and the outcome and then performs clustering using only the features with the largest univariate scores. The difference between semi-supervised RPMM and supervised clustering is that semi-supervised RPMM applies the RPMM algorithm of [Houseman and others \(2008\)](#) to the surviving features rather than a more conventional k-means or hierarchical clustering model.

These methods have been successfully identified clinically relevant subtypes of cancer in many different studies ([Bair and Tibshirani, 2004](#); [Bullinger and others, 2004](#); [Chinnaiyan and others, 2008](#); [Koestler and others, 2010](#)). However, these methods have significant limitations. In particular, both supervised clustering and semi-supervised RPMM require a user to choose the number of features that are used to form the clusters, and the results of these methods can depend heavily on the number of “significant” features selected. Moreover, it is very unlikely that these methods will successfully identify the truly significant features that define the clusters while excluding irrelevant features. We propose an alternative semi-supervised clustering method that is applicable in situations where the outcome variable is a noisy surrogate for an underlying biological subtype of interest. It is based on a modification of the “sparse clustering” algorithm of

Witten and Tibshirani (2010) which we call preweighted sparse clustering. We will show that our proposed method produces more accurate results than competing methods in a several simulated data sets as well as real-world data studies of chronic pain and cancer.

2. METHODS

2.1 *Preweighted Sparse Clustering*

Suppose that we wish to cluster the $n \times p$ data matrix X , where n is the number of observations and p is the number of features. Assume that the clusters only differ with respect to some subset of the features. Witten and Tibshirani (2010) propose a method that called “sparse clustering” to solve this problem. A brief description of the sparse clustering is as follows: Let $d_{i,i',j}$ be any dissimilarity measure between observations i and i' with respect to feature j . (Throughout the remainder of this discussion, we will assume that $d_{i,i',j} = (X_{ij} - X_{i'j})^2$ the Euclidean distance between X_{ij} and $X_{i'j}$.) Then Witten and Tibshirani (2010) propose to identify clusters C_1, C_2, \dots, C_K and weights w_1, w_2, \dots, w_p that maximize the weighted between-cluster sum of squares

$$\sum_{j=1}^p w_j \left(\frac{1}{n} \sum_{i=1}^n \sum_{i'=1}^n d_{i,i',j} - \sum_{k=1}^K \frac{1}{n_k} \sum_{i,i' \in C_k} d_{i,i',j} \right) \quad (2.1)$$

subject to the constraints $\sum_j w_j^2 = 1$, $\sum_j |w_j| < s$, and $w_j \geq 0$ for all j , where s is a tuning parameter and n_k is the number of elements in cluster k . To maximize (2.1), Witten and Tibshirani (2010) use the following algorithm:

1. Initialize the weights as $w_1 = w_2 = \dots = w_p = 1/\sqrt{p}$.
2. Fix the w_i 's and identify C_1, C_2, \dots, C_K to maximize (2.1). This can be done by applying the standard k-means clustering method to the $n \times n$ dissimilarity matrix where the (i, i') element is $\sum_j w_j d_{i,i',j}$.
3. Fix the C_i 's and identify w_1, w_2, \dots, w_p to maximize (2.1) subject to the constraints that

$\sum_j w_j^2 = 1$ and $\sum_j |w_j| < s$. See [Witten and Tibshirani \(2010\)](#) for a description of how the optimal w_i 's are calculated.

4. Repeat steps 2 and 3 until convergence.

This procedure requires a user to choose the number of clusters k and the tuning parameter s . We will not discuss methods for choosing these parameters; see [Witten and Tibshirani \(2010\)](#) for an algorithm for choosing s , and see [Tibshirani and others \(2001\)](#), [Sugar and James \(2003\)](#), or [Tibshirani and Walther \(2005\)](#) for several possible methods for choosing k .

Although this method produces impressive results in a wide variety of problems, it tends to identify clusters that are dominated by highly correlated features with high variance, which may not be interesting biologically. It also does not consider the values of any outcome variables that may exist. Thus, in the situation illustrated in Figure 1, there is no guarantee that the clusters identified by this method will be associated with the outcome of interest. To overcome this issue, we propose the following simple modification of sparse clustering, which we call preweighted sparse clustering. The preweighted sparse clustering algorithm is described below:

1. Run the sparse clustering algorithm, as described above.
2. For each feature, calculate the F-statistic, F_i , (and associated p-value p_i) for testing the null hypothesis that the mean value of the feature i does not vary across the clusters.
3. For each feature i , define:

$$w_i = \begin{cases} 1/\sqrt{m} & \text{if } p_i \geq \alpha \\ 0 & \text{otherwise} \end{cases}$$

where m is the number of p_i 's such $p_i \geq \alpha$.

4. Run the sparse clustering algorithm using these w_i 's (beginning with step 2) and continuing until convergence.

In other words, the preweighted sparse clustering algorithm first performs conventional sparse clustering. It then identifies features whose mean values differ across the clusters. Then the sparse clustering algorithm is run a second time, but rather than giving equal weights to all features in the first step, this preweighted version of sparse clustering assigns a weight of 0 to all features that differed across the first set of clusters. The motivation is that this procedure will identify secondary clusters that would otherwise be obscured by clusters which have a larger dissimilarity measure (such as the situation illustrated in Figure 1). [Nowak and Tibshirani \(2008\)](#) refer to this type of procedure as “complementary clustering.”

This procedure requires one to choose a p-value threshold α for deciding which features should be give nonzero weight. An obvious choice is $\alpha = 0.05/p$, where p is the number of features. However, the user may choose a less or more stringent cutoff depending on the sample size and other considerations. Also note that this procedure may be repeated multiple times if the secondary clusters identified are still unrelated to the biological outcome of interest.

2.2 *Supervised Sparse Clustering*

The preweighted sparse clustering algorithm described above is an unsupervised method, since it does not require or use any kind of outcome variable. If an outcome variable is available and the objective is to identify clusters associated with the outcome variable, one may use the following variant of preweighted sparse clustering to incorporate such data, which we call supervised sparse cluster. The supervised sparse clustering procedure is described below:

1. Let T_i be a measure of the strength of the association between the i th feature and the outcome variable. (If the outcome variable is dichotomous, T_i could be a t-statistic, or if the outcome variable is a survival time, T_i could be a univariate Cox score.) Let $T_{(1)}, T_{(2)}, \dots, T_{(p)}$ denote the order statistics of the T_i 's.

2. Run the sparse clustering algorithm with initial weights w_1, w_2, \dots, w_p , where

$$w_i = \begin{cases} 1/\sqrt{m} & \text{if } |T_i| \geq |T_{(p-m+1)}| \\ 0 & \text{otherwise} \end{cases}$$

3. Repeat steps 2 and 3 from the standard sparse clustering algorithm until convergence.

In other words, supervised sparse clustering chooses the initial weights for the sparse clustering algorithm by giving nonzero weights to the features that are most strongly associated with the outcome variable. Note that that no initial clustering step is required. This is similar to the semi-supervised clustering method of [Bair and Tibshirani \(2004\)](#) and the semi-supervised RPMM method of [Koestler and others \(2010\)](#).

The supervised sparse clustering procedure requires the choice of a tuning parameter m , which is the number of features to be given nonzero weight in the first step. Our experience suggests that the procedure tends to give very similar results for a wide variety of different values of m ; therefore, optimizing the procedure with respect to this tuning parameter is unnecessary. As a default we suggest $m = \sqrt{p}$, where p is the number of features. We will use this default throughout this manuscript unless otherwise noted.

2.3 Simulated Data Sets

To test the preweighted sparse clustering algorithm, we generated a series of 10 simulated 5000×200 data sets as follows:

$$X_{ij} = \begin{cases} 1 + \epsilon_{ij} & \text{if } i \leq 50, j \leq 100 \\ -1 + \epsilon_{ij} & \text{if } i \leq 50, j > 100 \\ 2 + \epsilon_{ij} & \text{if } 50 < i \leq 250, 50 < j \leq 150 \\ -2 + \epsilon_{ij} & \text{if } 50 < i \leq 250, j \leq 50 \text{ or } j > 150 \\ \epsilon_{ij} & \text{if } i > 250 \end{cases}$$

Here the ϵ_{ij} 's are iid standard normal random variables. In this simulation, each row represents a feature and each column represents an observation, as is conventional in microarray data sets. Note that this is essentially the same situation depicted in [Figure 1](#). Features 51-250 form one

set of clusters and features 1-50 form a secondary set of clusters.

The objective of this simulation is to determine if preweighted clustering can identify both sets of clusters in this data set. For each simulated data set, we applied the preweighted sparse clustering algorithm as described in Section 2.1 and obtained the clusters produced by the initial sparse clustering and the secondary clusters produced after preweighting. Both sets of cluster assignments were compared to the true clusters.

We also generated a series of 10 simulated data sets to test the supervised sparse clustering algorithm. Specifically, we generated 10 5000×200 data matrices X where

$$X_{ij} = \begin{cases} 1 + \epsilon_{ij} & \text{if } i \leq 50, j \leq 100 \\ 2 + \epsilon_{ij} & \text{if } i \leq 50, j > 100 \\ 2I(u_{ij} < 0.4) + \epsilon_{ij} & \text{if } 51 \leq i \leq 100 \\ 0.5I(u_{ij} < 0.7) + \epsilon_{ij} & \text{if } 101 \leq i \leq 200 \\ 1.5I(u_{ij} < 0.3) + \epsilon_{ij} & \text{if } 201 \leq i \leq 300 \\ \epsilon_{ij} & \text{if } i > 300 \end{cases}$$

Here $I(x)$ is an indicator function, and the u_{ij} 's are iid uniform random variables on $(0, 1)$. The ϵ_{ij} 's are iid standard normal, as before. We also defined the binary outcome variable y as follows:

$$y_i = \begin{cases} 0 + I(u_i < 0.3) & \text{if } 1 \leq i \leq 100 \\ 1 - I(u_i < 0.3) & \text{if } 101 \leq i \leq 200 \end{cases}$$

(In the above, once again $I(x)$ is an indicator function and the u_i 's are iid uniform random variables on $(0, 1)$.) This simulation is similar to the scenario illustrated in Figure 1. We assume that the first 50 features are the biologically relevant features of interest. In other words, a clustering algorithm that achieves perfect accuracy should assign observations 1-100 to one cluster and observations 101-200 to a separate cluster. Features 51-100, 101-200, and 201-300 also form clusters, but these clusters are not related to the biological outcome of interest. The outcome variable y also observed is a “noisy surrogate” for the true clusters. This y is related to the true clusters, but 30% of the y_i 's are misclassified. This is consistent with what we might expect to observe in a study of chronic pain, where the only observed outcome variable is a patient's subjective pain report that is not always a reliable indicator of case status.

The objective of this simulation is to determine if supervised sparse clustering can correctly identify the clusters that are associated with the y_i 's (as opposed to the other sets of spurious clusters). Supervised sparse clustering was applied to each of the 10 simulated data sets. Three other methods were also considered, namely conventional sparse clustering, the semi-supervised clustering method of [Bair and Tibshirani \(2004\)](#), and conventional 2-means clustering on the first three principal components of the data set. We also attempted to apply the semi-supervised RPMM method of [Koestler and others \(2010\)](#) to these simulated data sets, but in each case the procedure returned a singleton cluster.

2.4 OPPERA Data

We applied our preweighted sparse clustering method to a data set collected in the Orofacial Pain: Prospective Evaluation and Risk Assessment (OPPERA) study. OPPERA is a prospective cohort study to identify risk factors for temporomandibular disorders (TMD). OPPERA recruited a total of 3443 study subjects at four U.S. study sites from May 2006 to November 2008. The cohort included 185 subjects with chronic TMD and 3258 healthy controls. The initially TMD-free individuals completed a quarterly questionnaire assessing TMD pain symptoms, and those reporting symptoms were invited for a follow up exam to determine if they had developed first-onset TMD. The median follow up period was 2.8 years, and a total of 260 participants developed TMD over the course of the study. For a more detailed description of the OPPERA study, see [Maixner and others \(2011\)](#) or [Slade and others \(2011\)](#).

We applied our preweighted sparse clustering algorithm to the data from the OPPERA baseline case-control study, which includes half (i.e. 1632) of the initially TMD-free individuals and all 185 chronic TMD cases. See [Slade and others \(2011\)](#) for a description of how this cohort was selected. We included all of the measures of experimental pain sensitivity, psychological distress, and autonomic function that were previously analyzed in OPPERA. See [Greenspan and others](#)

(2011), [Fillingim *and others* \(2011\)](#), and [Maixner *and others* \(2011\)](#) for a more detailed description of these variables. A total of 116 predictor variables were used, including 33 measures of experimental pain sensitivity, 39 measures of psychological distress, and 44 measures of autonomic function.

In our initial analysis of the OPPERA data, we applied the preweighted sparse clustering algorithm as outlined in Section 2.1. Conventional sparse 2-means clustering was applied to the data set, after which the features that showed strongest mean differences across the clusters were given a weight of 0 when the preweighted version of sparse clustering was applied. All features were normalized to have mean 0 and standard deviation 1 prior to performing the clustering. The association between both the primary clusters and secondary clusters and chronic TMD was evaluated by calculating odds ratios and performing a chi-square test of the null hypothesis of no association between the clusters and TMD case status. We also applied our supervised sparse clustering method to the OPPERA data, as well as the other three clustering methods considered in the simulation study in Section 2.3 (namely sparse clustering, semi-supervised clustering, and clustering on the principal component scores). The association between the clusters produced by each of these methods and chronic TMD was again evaluated by calculating the odds ratios.

2.5 *Breast Cancer Microarray Data*

We also applied our supervised sparse clustering algorithm to the breast cancer microarray data of [van 't Veer *and others* \(2002\)](#). This data set includes data for 78 subjects with lymph-node-negative breast cancer. Gene expression data for 4750 genes are recorded for each subject, as well as survival times and outcomes. Survival times ranged from 0.27 to 13.4 years, with an average time of 5.99 years. The objective was to identify genetic subtypes (i.e. clusters) using the gene expression data that could be used to predict the prognosis of breast cancer patients.

We applied our supervised sparse 2-means clustering method to this data set as well con-

ventional sparse clustering, semi-supervised clustering, and clustering on the PCA scores. Before applying any of the clustering methods, the data were randomly partitioned in a training set and a test set, each of which consisted of 39 observations. Each clustering method was applied to the training data. To identify the “most significant” genes prior to applying supervised sparse clustering and semi-supervised clustering, the association between each gene and survival was evaluated by calculating the univariate Cox score for each gene. See [Beer *and others* \(2002\)](#) or [Bair and Tibshirani \(2004\)](#) for more information. For each set of clusters, a nearest shrunken centroid model ([Tibshirani *and others*, 2002](#)) was fit to the clusters in the training data and then applied to the test data to predict cluster assignments on the test data. The association between the predicted clusters in the test set and survival was evaluated using log-rank tests for each clustering method.

3. RESULTS

3.1 *Simulated Data Sets*

For the first simulation scenario, preweighted sparse clustering performed exactly as expected. For all 10 simulated data sets, the preliminary supervised clustering step assigned observations 51-150 to one cluster and observations 1-50 and 151-200 to the second cluster. After applying the preweighting, the procedure assigned observations 1-100 to one cluster and features 101-200 to the second cluster. Thus, preweighted sparse clustering successfully identifies both sets of clusters in this simulation scenario.

To evaluate the results of each method in the second simulation scenario, for each set of predicted clusters we identified the cluster that contained 51 or more of observations 1-100. This cluster was called cluster 1 and the other cluster was called cluster 2. Note that if the clustering method correctly identifies the clusters of interest, then observations 1-100 should be assigned to cluster 1 and observations 101-200 should be assigned to cluster 2. Thus, for each set of clusters,

we calculated the number of misclassified observations by counting the number of observations 101-200 that were assigned to cluster 1 plus the number of observations 1-100 that were assigned to cluster 2.

The mean number of misclassified observations (and associated standard errors) for each method are shown in Table 1. Supervised sparse clustering produced the lowest error rate of all the methods considered. (Indeed, supervised sparse clustering produced no more than 1 misclassified observation for 9 of the 10 simulations.) Semi-supervised cluster occasionally identified the correct clusters, but produced unsatisfactory results in several of the simulations. Conventional sparse clustering and 2-means clustering on the principal component scores produced poor results in all the simulated data sets.

3.2 *OPPERA Data*

We applied the preweighted sparse 2-means clustering method to the *OPPERA* data. The weights for both the primary and secondary clusters are shown in Figure 3. Observe that the measures of autonomic function had the largest feature weights for the primary clusters, whereas the measures of psychological distress had the largest feature weights for the secondary clusters. Thus, the preweighted sparse clustering method revealed a biologically meaningful set of secondary clusters that was not identified by the conventional sparse clustering algorithm. The association between both the primary and secondary clusters and chronic TMD is shown in Table 2. Observe that there is a significantly higher proportion of TMD cases in cluster 2 for both the primary and secondary clusters, but the association is stronger in the secondary clusters ($OR = 1.93, p = 3.2 \times 10^{-5}$) than in the primary clusters ($OR = 1.64, p = 0.002$). Thus, if the primary objective is to identify clusters associated with TMD, the secondary clusters are preferable to the primary clusters, indicating that the secondary clusters are not only biologically meaningful but may be more clinically relevant than the primary clusters.

We also applied supervised sparse 3-means clustering (as well as three other variants of 3-means clustering discussed earlier) to the OPPERA data. The results are shown in Table 3. While all four methods identified clusters that were associated with TMD, the clusters produced by supervised sparse clustering and semi-supervised clustering were much more strongly associated with TMD than the clusters produced by the methods that did not consider an outcome variable. This suggests that clustering methods that consider an outcome variable may do a better job of identifying biologically relevant clusters than methods that do not consider this information.

3.3 Breast Cancer Microarray Data

For each clustering method, the chi-square statistic (and associated p-value) for the log-rank test of the null hypothesis of no difference in survival for the test set clusters are shown in Table 4. The results for conventional sparse clustering are not reported in this table, since the sparse clustering algorithm grouped all but one training observation into the same cluster, and as a result all 39 test set observations were predicted to belong to the same cluster. We observed that supervised sparse clustering identifies test set clusters that are significantly associated with survival (at the $p < 0.05$ level) whereas the other two methods considered do not. This indicates that supervised sparse clustering can identify biologically meaningful and clinically relevant subtypes of cancer based on microarray data. The fact that the predicted clusters were associated with survival on an independent test set suggests that this finding is not merely the result of overfitting.

4. DISCUSSION

Cluster analysis is frequently used to identify subtypes in complex data sets. In many cases, the primary objective of the cluster analysis is to identify clusters that offer new insight into a biological question of interest or that can be used to more precisely phenotype (and hence diagnose and treat) a particular disease. However, in many cases, the clusters identified by conventional

clustering methods are dominated by a subset of the features that are not interesting biologically or clinically. Despite the fact that this problem is very common in cluster analysis, relatively few methods have been proposed to overcome this problem.

The idea of “complementary clustering,” that is, the identification of secondary clusters in a data set after removing the effects of a primary cluster of lesser interest, was first proposed by [Nowak and Tibshirani \(2008\)](#). [Witten and Tibshirani \(2010\)](#) propose an alternative method for complementary clustering (which they called “complementary sparse clustering”) based on their sparse clustering algorithm. Both methods are innovative and useful. However, their main drawback is that they can only be used with hierarchical clustering. To our knowledge, our proposed method is the first complementary clustering algorithm that may be applied to k-means clustering or other clustering methods. Although we have only considered preweighted k-means clustering in this study, our methodology is easily applicable to sparse hierarchical clustering or any other clustering method that can be used within the sparse clustering framework of [Witten and Tibshirani \(2010\)](#).

The problem of finding clusters that are associated with an outcome variable has also not been studied extensively. Previously proposed methods include the semi-supervised clustering method of [Bair and Tibshirani \(2004\)](#) and the semi-supervised RPMM method of [Koestler and others \(2010\)](#). Semi-supervised clustering produces useful results in a variety of circumstances, but the clusters produced by semi-supervised clustering can vary depending on the choice of tuning parameters and sometimes have poor reproducibility. Semi-supervised clustering can also fail to identify the true clusters of interest when the association between these clusters and the observed outcome is noisy, as we saw in [Section 3.1](#). Likewise, a drawback of semi-supervised RPMM is that it can fail to detect that clusters exist in a data set. (Indeed, semi-supervised RPMM produced a singleton cluster in each of the examples we considered in the present study.) Supervised sparse clustering has been shown to overcome these shortcomings and can produce reproducible clusters

associated with an outcome in situations where semi-supervised cluster fails (see Section 3.3).

One shortcoming of the proposed preweighted sparse clustering is the fact that the clusters obtained may vary with respect to the choice of the tuning parameter s in the sparse clustering algorithm (see Section). The question of how to choose this tuning parameter has not been studied extensively. Witten and Tibshirani (2010) propose a method for choosing s based on permuting the columns of the data, but in our experience this method tends to produce values of s that are too large, which sometimes results in clusters that are not associated (or less strongly associated) with the outcome of interest. Choosing a smaller value of s may produce better results. The question of how to choose this tuning parameter is an area for further study.

Despite this limitation, we believe that preweighted sparse clustering and supervised sparse clustering are powerful tools for solving an understudied problem. These methods can be used to identify biologically meaningful clusters in data sets that may not be detected by existing methods. More importantly, these methods can be used to identify clinically relevant subtypes of diseases like TMD and cancer, ultimately leading to better treatment options.

5. SOFTWARE

Software in the form of R code, together with a sample input data set and complete documentation is available on request from the corresponding author (eaheron@tcd.ie).

6. SUPPLEMENTARY MATERIAL

Supplementary material is available online at <http://biostatistics.oxfordjournals.org>.

ACKNOWLEDGMENTS

Conflict of Interest: None declared.

REFERENCES

- BAIR, ERIC, HASTIE, TREVOR, PAUL, DEBASHIS AND TIBSHIRANI, ROBERT. (2006). Prediction by supervised principal components. *Journal of the American Statistical Association* **101**(473), 119–137.
- BAIR, ERIC AND TIBSHIRANI, ROBERT. (2004, 04). Semi-supervised methods to predict patient survival from gene expression data. *PLoS Biol* **2**(4), e108.
- BASU, SUGATO, BANERJEE, ARINDAM AND MOONEY, RAYMOND J. (2002). Semi-supervised clustering by seeding. In: *Proceedings of the Nineteenth International Conference on Machine Learning*, ICML '02. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc. pp. 27–34.
- BEER, DAVID G., KARDIA, SHARON L.R., HUANG, CHIANG-CHING, GIORDANO, THOMAS J., LEVIN, ALBERT M., MISEK, DAVID E., LIN, LIN, CHEN, GUOAN, GHARIB, TAREK G., THOMAS, DAFYDD G., LIZYNESS, MICHELLE L., KUICK, RORK, HAYASAKA, SATORU, TAYLOR, JEREMY M.G., IANNETTONI, MARK D., ORRINGER, MARK B. *and others.* (2002). Gene-expression profiles predict survival of patients with lung adenocarcinoma. *Nature Medicine* **8**, 816–824.
- BHATTACHARJEE, ARINDAM, RICHARDS, WILLIAM G., STAUNTON, JANE, LI, CHENG, MONTI, STEFANO, VASA, PRIYA, LADD, CHRISTINE, BEHESHTI, JAVAD, BUENO, RAPHAEL, GILLETTE, MICHAEL, LODA, MASSIMO, WEBER, GRIFFIN, MARK, EUGENE J., LANDER, ERIC S., WONG, WING, JOHNSON, BRUCE E., GOLUB, TODD R., SUGARBAKER, DAVID J. *and others.* (2001). Classification of human lung carcinomas by mrna expression profiling reveals distinct adenocarcinoma subclasses. *Proceedings of the National Academy of Sciences* **98**, 13790–13795.
- BRUEHL, STEPHEN, HARDEN, R.NORMAN, GALER, BRADLEY S., SALTZ, SAMUEL, BACKONJA,

- MIROSLAV AND STANTON-HICKS, MICHAEL. (2002). Complex regional pain syndrome: are there distinct subtypes and sequential stages of the syndrome? *Pain* **95**(12), 119 – 124.
- BULLINGER, LARS, DÖHNER, KONSTANZE, BAIR, ERIC, FRÖHLING, STEFAN, SCHLENK, RICHARD, TIBSHIRANI, ROBERT, DÖHNER, HARTMUT AND POLLACK, JONATHAN R. (2004). Gene expression profiling identifies new subclasses and improves outcome prediction in adult myeloid leukemia. *The New England Journal of Medicine* **350**, 1605–1616.
- CHINNAIYAN, ARUL M., LIPPMAN, MARC E., YU, JIANJUN, YU, JINDAN, CORDERO, KEVIN E., JOHNSON, MICHAEL D., GHOSH, DEBASHIS AND RAE, JAMES M. (2008, JAN). A transcriptional fingerprint of estrogen in human breast cancer predicts patient survival. *NEOPLASIA* **10**(1), 79–88.
- DAVIS, PHILIP J., REEVES, JOHN L., GRAFF-RADFORD, STEVEN B., HASTIE, BARBARA A. AND NALIBOFF, BRUCE D. (2003). Multidimensional subgroups in migraine: Differential treatment outcome to a pain medicine program. *Pain Medicine* **4**(3), 215–222.
- DIATCHENKO, LUDA, NACKLEY, ANDREA G., SLADE, GARY D., FILLINGIM, ROGER B. AND MAIXNER, WILLIAM. (2006). Idiopathic pain disorders pathways of vulnerability. *PAIN* **123**(3), 226 – 230.
- DWORKIN, S.F. AND LERESCHE, L. (1992). Research diagnostic criteria for temporomandibular disorders: review, criteria, examinations and specifications, critique. *Journal of Craniomandibular Disorders* **6**(4), 301–355.
- FILLINGIM, R. B., SLADE, G. D., DIATCHENKO, L., DUBNER, R., GREENSPAN, J. D., KNOTT, C., OHRBACH, R. AND MAIXNER, W. (2011, Nov). Summary of findings from the OPPERA baseline case-control study: implications and future directions. *J Pain* **12**(11 Suppl), T102–107.
- FRIEDMAN, JEROME H. AND MEULMAN, JACQUELINE J. (2004). Clustering objects on subsets

- of attributes (with discussion). *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **66**(4), 815–849.
- GHOSH, DEBASHIS AND CHINNAIYAN, ARUL M. (2002). Mixture modelling of gene expression data from microarray experiments. *Bioinformatics* **18**(2), 275–286.
- GREENSPAN, J. D., SLADE, G. D., BAIR, E., DUBNER, R., FILLINGIM, R. B., OHRBACH, R., KNOTT, C., MULKEY, F., ROTHWELL, R. AND MAIXNER, W. (2011, Nov). Pain sensitivity risk factors for chronic TMD: descriptive data and empirically identified domains from the OPPERA case control study. *J Pain* **12**(11 Suppl), 61–74.
- HANDL, JULIA AND KNOWLES, JOSHUA. (2006). On semi-supervised clustering via multiobjective optimization. In: *Proceedings of the 8th annual conference on Genetic and evolutionary computation*, GECCO '06. New York, NY, USA: ACM. pp. 1465–1472.
- HASTIE, BARBARA A., III, JOSEPH L. RILEY, ROBINSON, MICHAEL E., GLOVER, TONI, CAMPBELL, CLAUDIA M., STAUD, ROLAND AND FILLINGIM, ROGER B. (2005). Cluster analysis of multiple experimental pain modalities. *Pain* **116**(3), 227 – 237.
- HOUSEMAN, E ANDRES, CHRISTENSEN, BROCK, YEH, RU-FANG, MARSIT, CARMEN, KARAGAS, MARGARET, WRENSCH, MARGARET, NELSON, HEATHER, WIEMELS, JOSEPH, ZHENG, SHICHUN, WIENCKE, JOHN *and others*. (2008). Model-based clustering of dna methylation array data: a recursive-partitioning algorithm for high-dimensional data arising as a mixture of beta distributions. *BMC Bioinformatics* **9**(1), 365.
- JAMISON, ROBERT N., ROCK, DANIEL L. AND PARRIS, WINSTON C. V. (1988). Empirically derived symptom checklist 90 subgroups of chronic pain patients: A cluster analysis. *Journal of Behavioral Medicine* **11**, 147–158. 10.1007/BF00848262.
- KOESTLER, DEVIN C., MARSIT, CARMEN J., CHRISTENSEN, BROCK C., KARAGAS, MAR-

- GARET R., BUENO, RAPHAEL, SUGARBAKER, DAVID J., KELSEY, KARL T. AND HOUSEMAN, E. ANDRES. (2010). Semi-supervised recursively partitioned mixture models for identifying cancer subtypes. *Bioinformatics* **26**(20), 2578–2585.
- LAPOINTE, JACQUES, LI, CHUNDE, VAN DE RIJN, MATT, HUGGINS, JOHN P., BAIR, ERIC, MONTGOMERY, KELLI, FERRARI, MICHELLE, RAYFORD, WALTER, EKMAN, PETER, DEMARZO, ANGELO M., TIBSHIRANI, ROBERT, BOTSTEIN, DAVID, BROWN, PATRICK O., BROOKS, JAMES D. *and others*. (2004). Gene expression profiling identifies clinically relevant subtypes of prostate cancer. *Proceedings of the National Academy of Sciences* **101**, 811–816.
- MAIXNER, W., DIATCHENKO, L., DUBNER, R., FILLINGIM, R. B., GREENSPAN, J. D., KNOTT, C., OHRBACH, R., WEIR, B. AND SLADE, G. D. (2011a, Nov). Orofacial pain prospective evaluation and risk assessment study—the OPPERA study. *J Pain* **12**(11 Suppl), 4–11.
- MAIXNER, W., GREENSPAN, J. D., DUBNER, R., BAIR, E., MULKEY, F., MILLER, V., KNOTT, C., SLADE, G. D., OHRBACH, R., DIATCHENKO, L. *and others*. (2011b, Nov). Potential autonomic risk factors for chronic TMD: descriptive data and empirically identified domains from the OPPERA case-control study. *J Pain* **12**(11 Suppl), 75–91.
- NOWAK, GEN AND TIBSHIRANI, ROBERT. (2008). Complementary hierarchical clustering. *Biostatistics* **9**(3), 467–483.
- PAN, WEI AND SHEN, XIAOTONG. (2007, May). Penalized model-based clustering with application to variable selection. *J. Mach. Learn. Res.* **8**, 1145–1164.
- QU, YI AND XU, SHIZHONG. (2004). Supervised cluster analysis for microarray data based on multivariate gaussian mixture. *Bioinformatics* **20**(12), 1905–1913.
- RAFTERY, ADRIAN E AND DEAN, NEMA. (2006). Variable selection for model-based clustering. *Journal of the American Statistical Association* **101**(473), 168–178.

- ROSENWALD, ANDREAS, WRIGHT, GEORGE, CHAN, WING C., CONNORS, JOSEPH M., CAMPO, ELIAS, FISHER, RICHARD I., GASCOYNE, RANDY D., MULLER-HERMELINK, H. KONRAD, SMELAND, ERLAND B. AND STAUDT, LOUIS M. (2002). The use of molecular profiling to predict survival after chemotherapy for diffuse large b-cell lymphoma. *The New England Journal of Medicine* **346**, 1937–1947.
- SANFILIPPO, JOSEPH AND ERB, TERESA. (2008). Evaluation and management of dysmenorrhea in adolescents. *Clinical Obstetrics & Gynecology* **51**(2), 257–267.
- SLADE, G. D., BAIR, E., BY, K., MULKEY, F., BARAIAAN, C., ROTHWELL, R., REYNOLDS, M., MILLER, V., GONZALEZ, Y., GORDON, S., RIBEIRO-DASILVA, M., LIM, P. F., GREENSPAN, J. D., DUBNER, R., FILLINGIM, R. B., DIATCHENKO, L., MAIXNER, W., DAMPIER, D., KNOTT, C. *and others*. (2011, Nov). Study methods, recruitment, sociodemographic findings, and demographic representativeness in the OPPERA study. *J Pain* **12**(11 Suppl), 12–26.
- SORLIE, T., PEROU, CM., TIBSHIRANI, R., AAS, T., GEISLER, S., JOHNSEN, H., HASTIE, T., EISEN, MB., VAN DE RIJN, M., JEFFREY, SS., THORSEN, T., QUIST, H., MATESE, JC., BROWN, PO., BOTSTEIN, D., LONNING, PE. *and others*. (2001). Gene expression patterns of breast carcinomas distinguish tumor subclasses with clinical implications. *Proceedings of the National Academy of Sciences* **98**, 10969–74.
- SUGAR, CATHERINE A AND JAMES, GARETH M. (2003). Finding the number of clusters in a dataset. *Journal of the American Statistical Association* **98**(463), 750–763.
- TARI, LUIS, BARAL, CHITTA AND KIM, SEUNGCHAN. (2009). Fuzzy c-means clustering with prior biological knowledge. *Journal of Biomedical Informatics* **42**(1), 74 – 81.
- TIBSHIRANI, R., HASTIE, T., NARASIMHAN, B. AND CHU, G. (2002, May). Diagnosis of multiple cancer types by shrunken centroids of gene expression. *Proc. Natl. Acad. Sci. U.S.A.* **99**(10), 6567–6572.

- TIBSHIRANI, ROBERT AND WALTHER, GUENTHER. (2005). Cluster validation by prediction strength. *Journal of Computational and Graphical Statistics* **14**(3), 511–528.
- TIBSHIRANI, ROBERT, WALTHER, GUENTHER AND HASTIE, TREVOR. (2001). Estimating the number of clusters in a data set via the gap statistic. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **63**(2), 411–423.
- VAN 'T VEER, L. J., DAI, H., VAN DE VIJVER, M. J., HE, Y. D., HART, A. A., MAO, M., PETERSE, H. L., VAN DER KOOY, K., MARTON, M. J., WITTEVEEN, A. T., SCHREIBER, G. J., KERKHOVEN, R. M., ROBERTS, C., LINSLEY, P. S., BERNARDS, R. *and others*. (2002, Jan). Gene expression profiling predicts clinical outcome of breast cancer. *Nature* **415**(6871), 530–536.
- WITTEN, DANIELA M. AND TIBSHIRANI, ROBERT. (2010). A framework for feature selection in clustering. *Journal of the American Statistical Association* **105**(490), 713–726.

Table 1. *Results of the second simulation study. The following methods were applied to the simulated data set described in Section 2.3: 1) supervised sparse clustering, 2) sparse clustering, 3) semi-supervised clustering (Bair and Tibshirani, 2004), 4) 2-means clustering on the top 3 principal component (PCA) scores.*

	Sup. Sparse Clust.	Sparse Clust.	Semi-Sup. Clust.	Clust. on PCA
Mean	10	96.5	26.3	95.4
SE	9.9	0.8	10.6	1.1

Table 2. *Association between the primary and secondary clusters identified by preweighted sparse clustering and chronic TMD on the OPPERA data set. In each case, the cluster with the lower proportion of TMD cases was called cluster 1. Observe that there is a significantly higher proportion of TMD cases in cluster 2 for both the primary and secondary clusters, but the association is stronger in the secondary clusters ($OR = 1.93, p = 3.2 \times 10^{-5}$) than in the primary clusters ($OR = 1.64, p = 0.002$).*

	Clust. 1 (Primary)	Clust. 2 (Primary)	Clust. 1 (Secondary)	Clust. 2 (Secondary)
Controls	951	681	1133	499
Cases	85	100	100	85

Fig. 1. Artificial data set illustrating the shortcomings of conventional clustering methods. Suppose observations 1-100 are controls and observations 101-200 are cases. In this situation, one would be interested in the clusters formed by features 1-50, but most existing clustering methods would identify the clusters formed by features 51-250 (and hence fail to identify the cluster formed by features 1-50).

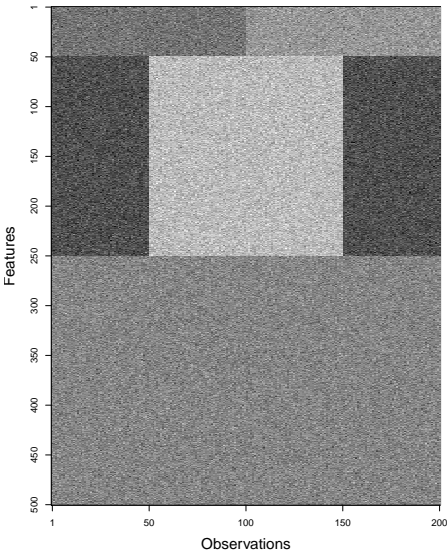


Fig. 2. Artificial example of a situation where the outcome variable is a “noisy surrogate” for the true clusters. In this artificial example, the density functions of the outcome variable for observations in each of two clusters are shown above. Observations in cluster 2 are more likely to have higher values of the outcome variable than observations in cluster 1, but there is considerable overlap between the two groups. Thus, classifying observations to clusters based solely on the outcome variable will result in a high misclassification error rate.

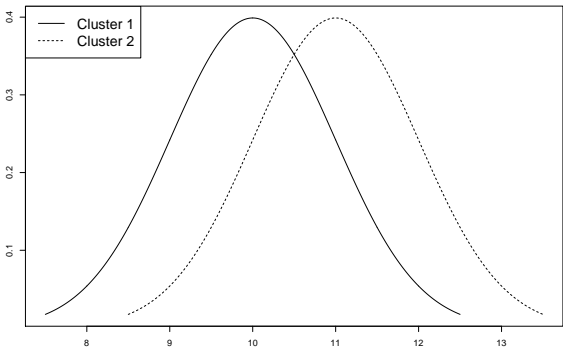


Fig. 3. Feature weights for the primary and secondary clusters identified by the preweighted sparse clustering method. In the figure below, features 1-33 are measures of experimental pain sensitivity, features 34-72 are measures of psychological distress, and features 73-116 are measures of autonomic function. We see that the primary clusters differ from one another primarily with respect to measures of autonomic functions and the secondary clusters differ primarily with respect to measures of psychological distress.

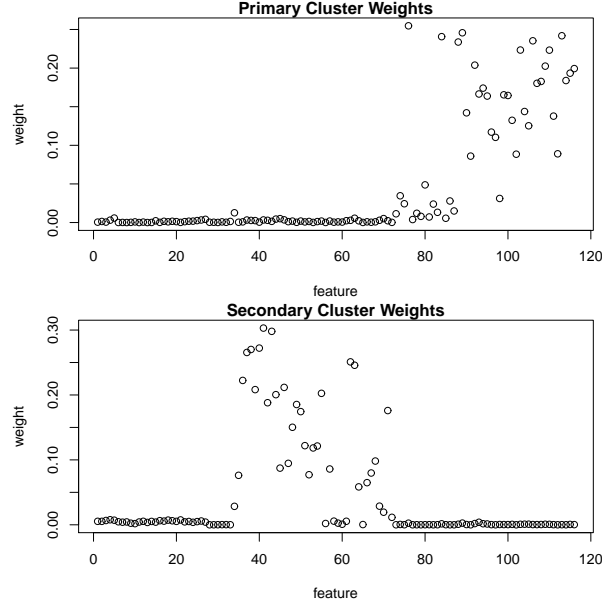


Table 3. Association between chronic TMD and the clusters produced by four different clustering methods on the OPPERA data. In each case, the cluster with the lowest proportion of TMD cases was called cluster 1 and the cluster with the highest proportion of TMD cases was called cluster 3. The odds ratio for TMD in each cluster (relative to cluster 1) was also calculated. The methods that use information about the outcome variable of interest (namely supervised sparse clustering and semi-supervised clustering) produce clusters more strongly associated with TMD than the other clustering methods.

		Cluster 1	Cluster 2	Cluster 3
Supervised Sparse Clustering	Controls	673	815	144
	Cases	20	131	34
	OR		5.4	7.9
Sparse Clustering	Controls	893	615	124
	Cases	78	77	30
	OR		1.4	2.8
Semi-Supervised Clustering	Controls	648	775	209
	Cases	13	81	91
	OR		5.2	21.7
Clustering on PCA Scores	Controls	756	608	268
	Cases	53	60	72
	OR		1.4	3.8

Table 4. *The association between the predicted clusters for the test data and survival for the breast cancer microarray data. For each method, a log-rank test was performed to test the null hypothesis of no difference in survival between the two predicted clusters. The chi-square statistic and associated p-values of each test is reported below. Note that the results for conventional sparse clustering are not shown in the table since all subjects in the test set were classified into the same cluster.*

	Chi-Square Statistic	P-value
Supervised Sparse Clustering	6.7	0.01
Semi-Supervised Clustering	2.9	0.09
Clustering on PCA Scores	2.9	0.09